

REMARKS

Claim rejections under 35 USC 103(a)

All the pending claims have been rejected under 35 USC 103(a) as being unpatentable over Kiick (2003/0200250) in view of Rowlands (2003/0200250), Fischer (6,438,672), Drottar (6,170,025), Agatsuma (7,237,099), and Chi (6,209,086). Claims 1, 7, 16, 21, and 25 are independent claims, from which the remaining claims ultimately depend. Applicant submits that the independent claims at least as currently amended are patentable, such that the remaining claims that are still pending and that have been rejected on this basis are patentable at least because they depend from independent base claims. Throughout the following discussion, claim 1 is discussed as representative of all the independent claims insofar as the rejection over Kiick in view of one or more other references in combination is concerned.

Applicant submits that the overall crux of these arguments is that the Examiner has not considered the claimed invention “as a whole” as is required – “the claimed invention as a whole must be considered” (MPEP sec. 2141.02.I.) – but rather has distilled the claimed invention down to its “gist,” which is improper. “Distilling an invention down to the ‘gist’ or ‘thrust’ of an invention disregards the requirement of analyzing the subject matter ‘as a whole.’” (MPEP sec 2141.02.II., citing W.L. Gore & Associates, Inc. v. Garlock, Inc., 721 F.2d 1540, 220 USPQ 303 (Fed. Cir. 1983)) All the claim limitations have to be considered in assessing the patentability of this claim over the prior art. “All words in a claim must be considered in judging the patentability of that claim against the prior art.” (Id., citing In re Wilson, 424 F.2d 1382, 1385, 165 USPQ 494, 496 (CCPA 1970)) Applicant respectfully requests that the Examiner keep this standard for patentability in mind when reviewing Applicant’s arguments below.

Applicant in particular now discusses two limitations of the invention that Applicant submits are not taught, disclosed, or suggested by the prior art in combination as relied upon by the Examiner.

*(1) First limitation not in the prior art*

Applicant notes that the claimed invention is limited to a very particular way in which interrupts for I/O devices are assigned among the nodes of a system. You have a first node to which the I/O device generating an interrupt is connected, and you have a second node at which the interrupt service routine for handling this interrupt is located. It is noted that the first node is "different" than the second node; i.e., they are not the same node. In the claimed invention, then, "if the first node . . . has a cache, memory, and at least one processor," then you assign the interrupt "to the first node." However, "if the first node does not have a cache, memory, and at least one processor, but the second node . . . does have a cache, memory, and at least one processor," then you assign the interrupt "to the second node." Furthermore, "if both the first node and the second node do not each have a cache, memory, and at least one processor," then you assign the interrupt "to a third node . . . having memory and at least one processor." It is noted that the third node is "different" than the first and the second nodes; i.e., they are not the same nodes.

To make this ordering even more clear, Applicant has amended the claimed invention as follows. Thus, "if the first node . . . has a cache, memory, and at least one processor," then you assign the interrupt to the first node "*even if the second node has a cache, memory, and at least one processor.*" Likewise, "if the first node does not have a cache, memory, and at least one processor, but the second node . . . does have a cache, memory, and least one processor," then you assign the interrupt to the second node "*even if a third node has a cache, memory, and at least one processor.*" In this way, the first node has priority over the second node as to interrupt assignment, and the second node has priority over any third node as to interrupt assignment, where the first, second, and third nodes each has a cache, memory, and at least one processor. For instance, if all three of the first, second, and third nodes have cache, memory, and at least one processor, then in the claimed invention, you always assign the interrupt to the first node instead

of to the second node and instead of to the third node. Likewise, if just the second and third nodes have cache, memory, and at least one processor, then in the claimed invention, you always assign the interrupt to the second node instead of to the third node.

In rejecting this aspect of the claimed invention, the Examiner has stated that paragraph 34 of Kiick “describes that interrupts should be assigned to the ‘closest’ processors,” and the Examiner “interprets this to mean the interrupts for the I/O devices *should be assigned to nodes to which they are connected or to nodes where the ISRs for the said I/O device reside.*” (Office action, p. 3.) The Examiner also notes that “it would be obvious to assign the interrupt to a node which has a processor and memory, since these elements are required to service an interrupt.” (Id., p. 4, relying upon Agatsuma, Drottar, and Fischer in particular.) On this basis it appears that the Examiner has concluded that the above-identified limitations of the claimed invention are taught, suggested, or disclosed by the prior art in combination.

Applicant disagrees, however, that these references in combinations rise to the level of disclosing, teaching, or suggesting the claim limitations in question. It is noted that there is a very particular order to which interrupt assignment is performed in the invention, as has been described above: if the (first) node to which the I/O device generating the interrupt has cache/memory/processors, then the interrupt is assigned to this node; if not, but if the (second) node having the interrupt service routine for this interrupt has cache/memory/processors, then the interrupt is assigned to this node; if not, then you assign the interrupt to a (third) node that does have memory/processors. Even in light of the teachings of the prior art ascribed by the Examiner, however, the prior art in combination does not teach, disclose, or suggest this particular order in assigning an interrupt.

For instance, the Examiner has stated that Kiick assigning interrupts to the closest processors means that they “should be assigned to nodes to which they are connected or to nodes where the ISRs for the said I/O device reside.” However, even if this is the case, this does not mean that the node to which the I/O device is connected has priority over the node at which the

ISR resides in receiving the interrupt, assuming both have cache/memory/processors. For example, consider the following scenario: the (first) node to which the I/O device is connected has cache/memory/processors, and the (second) node at which the ISR resides also has cache/memory/processors. In the claimed invention, there is a definitive selection of which node should have the interrupt assigned to it – the first node.

By comparison, in the prior art as combined by the Examiner, there is no definitive selection of which node should have the interrupt assigned to it. Kiick is unhelpful, because as interpreted by the Examiner, Kiick merely says that the interrupt should be assigned to either the first node or the second node. Agatsuma, Drottar, and Fischer are unhelpful, because as interpreted by the Examiner, they merely say you want to select a node to assign the interrupt to that has cache/memory/processors, and in our example, both the first and the second nodes have cache/memory/processors. Thus, against these teachings, disclosures, and suggestions of the prior art, one of ordinary skill within the art does not come away with the claimed invention. In other words, the prior art does not give preference to the first node in this example, where both the first node and the second node each has cache/memory/processors.

Stated another way, the teaching of Kiick that you want to assign an interrupt to the node that is “closest,” in combination with the teachings of Agatsuma, Drottar, and Fischer that you a node can only service an interrupt if it has cache/memory/processors, does not teach, disclose, or suggest the claimed invention. If both the node to which the I/O device is connected and the node having the ISR are “closest” per the interpretation of Kiick by the Examiner, and if both these nodes have cache/memory/processors such that they can service the interrupt per the interpretation of Agatsuma, Drottar, and Fischer by the Examiner, then the end result is that the prior art in combination seemingly results in the conclusion that *the prior art does not dictate to which of these two nodes should have the interrupt assigned thereto*. By comparison, the claimed invention does dictate to which of these two nodes the interrupt should be assigned: the node to

which the I/O device is connected always has the interrupt assigned to it unless this node does not have cache/memory/processors.

For this reason alone, Applicant respectfully submits that not all claim limitations are taught, suggested, and disclosed by the prior art in combination, rendering the claimed invention non-obvious and patentable over the prior art in combination.

*(2) Second limitation not in the prior art*

Applicant also notes that the claimed invention is limited to assigning interrupts in a round-robin manner, to dynamically modifying assignments of the interrupts based on actual performance characteristics, and for dynamically modifying assignments of interrupts that are performance critical. The Examiner has stated that the prior art in combination teaches, discloses, or suggests these limitations at least insofar as Kiick in paragraphs 25, 26, 28, and 31 teaches, discloses, or suggests these limitations. (Office action, p. 3.) Applicant respectfully disagrees.

Kiick does not actually discuss assigning interrupts and dynamically modifying interrupt assignments in these excerpts, such that the prior art in combination does not teach, disclose, or suggest these limitations. Rather, Kiick discusses assigning *interrupt service routines* (ISR's), and dynamically modifying assignments of ISR's. (See, e.g., paragraphs 25, 26, 28, and 31, which all discuss assignment and assignment modification of ISR's, where paragraph 24 notes that "processing of each I/O interrupt is facilitated by means of an interrupt service routine (ISR) that is assignable to a specific processor.") This is a key distinction.

The claimed invention, for instance, is limited to there being a second node that has an interrupt service routine, thus distinguishing between such an ISR and an interrupt, which is described in the claim language as being assigned, or having its assignment dynamically modified. Furthermore, the prior art itself distinguishes between an ISR and an interrupt, where paragraph 24 of Kiick notes that, as has been alluded to above, "processing of each I/O interrupt is facilitated by means of an interrupt service routine (ISR) that is assignable to a specific

processor.” Therefore, one cannot interpret the assignment of interrupts and the dynamic modification of interrupt assignments as in the claimed invention as being the same thing as the assignment of interrupt service routines and the dynamic modification of such interrupt service routines. Both the claimed invention and the prior art distinguish between an interrupt itself and the interrupt service routine (ISR) that services the interrupt.

As such, the limitations of the claimed invention relating to assigning interrupts and dynamically modifying interrupt assignments are not taught, disclosed, or suggested by the prior art in combination. The relied upon prior art does not actually disclose, teach or suggest assigning interrupts and dynamically modifying interrupt assignment, but rather discloses, teaches, and suggests assigning interrupt service routines and dynamically modifying interrupt service routine assignments. This is different. For at least this reason as well, then, the claimed invention is patentable over the prior art in combination.

Conclusion

Applicants have made a diligent effort to place the pending claims in condition for allowance, and request that they so be allowed. However, should there remain unresolved issues that require adverse action, it is respectfully requested that the Examiner telephone Mike Dryja, Applicant's representative, at 425-427-5094, so that such issues may be resolved as expeditiously as possible. For these reasons, and in view of the above amendments, this application is now considered to be in condition for allowance and such action is earnestly solicited.

Respectfully Submitted,



April 8, 2008

Date

Michael A. Dryja, Reg. No. 39,662  
Attorney/Agent for Applicant(s)

Law Offices of Michael Dryja  
1474 N Cooper Rd #105-248  
Gilbert, AZ 85233  
tel: 425-427-5094  
fax: 425-563-2098